

NETWORK COMMUNICATION AND MESSAGE
PROTOCOL FOR A MEDICAL PERFUSION SYSTEM

[0001] This application is a Continuation-in-Part Application of U.S. Application No. 09/030,989 filed on 26 February 1998, which is a Continuation of 5 U.S. Application No. 08/723,504 filed on 30 September 1996.

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0002] The present invention is directed to a medical perfusion system adapted to handle the selective oxygenation, filtering and recirculation of blood in 10 connection with various medical procedures.

[0003] A conventional perfusion system may be used to oxygenate, filter, and/or recirculate the blood of a patient during a medical procedure. Such a perfusion system may have a fluid conduit that removes blood from the patient during the medical procedure, a separate fluid conduit that returns blood to the 15 patient, one or more blood pumps that pump blood through the conduits, and a plurality of sensing devices, such as flow sensors and/or level sensors associated with blood pumps. The perfusion system may also include air embolus sensors, temperature sensors, flow occluders, etc.

[0004] Typically, a perfusion system is provided with a 20 configuration specifically designed to be used for a particular purpose. For example, one perfusion system may be specifically designed as a full-function heart/lung machine, while another perfusion system may be specifically designed as a ventricular-assist system. Although it may be possible to convert a perfusion system designed for one purpose to a perfusion system usable for a different 25 purpose, such reconfiguration is generally difficult and/or time-consuming.

SUMMARY OF THE INVENTION

[0005] The present invention involves a network communication and messaging protocol for use in a medical perfusion system. The protocol provides rules which govern how, among other things, information and data is conveyed

5 between the various devices which are connect to the network (e.g., controller and perfusion devices, including the adaptor pods associated with the perfusion devices). The rules define the different types of messages that are used in conveying information and data, as well as the formatting for each of those message types. The various messages provide the ability to configure the

10 perfusion system devices, establish links between the perfusion system devices, and convey various types of information and data, even when the main network controller is unavailable.

[0006] Accordingly, it is an object of the present invention to provide a perfusion system that can efficiently and effectively configure the various devices that make up the perfusion system.

[0007] It is another object of the present invention to provide a perfusion system in which links, such as feedback links and trigger links, between the various devices can be efficiently and effectively established.

[0008] It is still another object of the present invention to provide a mechanism for prioritizing the various messages that are used to communicate information and data between the devices that make up the perfusion system.

[0009] It is yet another object of the present invention to provide for device to device communication, even if the main network controller is unavailable.

[0010] In accordance with a first embodiment of the present invention, the above identified and other objectives are achieved by a medical perfusion system that includes a plurality of perfusion devices and a communications bus connecting each of the perfusion devices. The system also includes means for broadcasting a message to or from one of the perfusion devices, where the message contains a

data portion which identifies the message as belonging to one of a number of predefined message types.

[0011] In accordance with another embodiment of the present invention, the above identified and other objectives are achieved by a medical perfusion system that includes a plurality of perfusion devices, where one of the perfusion devices is a sensing device and a second one of the perfusion devices is a responding device. The system also includes a communications bus which connects each of the perfusion devices, where the sensing device has associated therewith means for broadcasting a message having a data portion which contains feedback data for controlling the operation of the responding device.

[0012] In accordance with still another embodiment of the present invention, the above identified and other objectives are achieved by a medical perfusion system that includes a plurality of perfusion devices, where one of the perfusion devices is a trigger source and a second one of the perfusion devices is a trigger respondent. The system also includes a communications bus which connects each of the perfusion devices, where the trigger source includes means for broadcasting a message having a data portion which alerts the trigger respondent to the existence of a particular condition.

[0013] In accordance with still another embodiment of the present invention, the above identified and other objectives are achieved by a medical perfusion system that includes a plurality of perfusion devices and a communications bus which connects each of the perfusion devices. The system also includes means for broadcasting a message to a first one of the perfusion devices, where the message includes a data portion containing configuration data associated with the first perfusion device.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] FIG. 1 is a block diagram of a preferred embodiment of a perfusion system in accordance with the invention;

- [0015] FIG. 2 is a perspective view of the main controller shown schematically in FIG. 1;
- [0016] FIG. 3 is a perspective view of one of the network extenders shown schematically in FIG. 1;
- 5 [0017] FIG. 4 is a perspective view of one of the adapter pods shown schematically in FIG. 1;
- [0018] FIGS. 5-7 illustrate a number of connector configurations;
- [0019] FIG. 8 is a perspective view of the main controller shown schematically in FIG. 1 with two network extenders and eight adapter pods
- 10 plugged therein;
- [0020] FIG. 9 is a block diagram of the main controller shown schematically in FIG. 1;
- [0021] FIG. 10 is a block diagram of one of the extender controllers shown schematically in FIG. 1;
- 15 [0022] FIG. 11 is a block diagram of one of the node controllers shown schematically in FIG. 1;
- [0023] FIG. 12 is a block diagram of one of the adapter pods shown schematically in FIG. 1;
- [0024] FIGS. 13A-13H are flowcharts illustrating the operation of the main controller shown in FIG. 1;
- 20 [0025] FIGS. 14A-14B are exemplary illustrations of a pair of perfusion circuit images generated on the display device of FIG. 9 during operation of the perfusion system;
- [0026] FIGS. 15A-15C are flowcharts illustrating the operation of the extender controllers shown in FIG. 1;
- 25 [0027] FIGS. 16A-16B are flowcharts illustrating the operation of the node controllers shown in FIG. 1;
- [0028] FIGS. 17A-17D are flowcharts illustrating the operation of the adapter pods shown in FIG. 1; and

[0029] FIG. 18 is a diagram of an exemplary digital data packet format.

DETAILED OF THE PREFERRED EMBODIMENTS

[0030] FIG. 1 illustrates a preferred embodiment of a medical perfusion system 10 in accordance with the invention. The perfusion system 10 is adapted to handle the selective oxygenation, filtering and recirculation of blood in connection with a number of different medical procedures. The perfusion system 10 may be placed in a number of different configurations, each of which corresponds to a different medical procedure. For example, the perfusion system 10 may be configured as a full-function heart/lung machine, a ventricular assist system, or a single-pump system that can be used for various purposes, such as to perform blood aspiration or myocardial protection during surgery.

[0031] Referring to FIG. 1, the main controller 20 is connected to a network extender 22a via a data/power bus 30a and to a network extender 22b via a data/power bus 30b. The network extender 22a includes an extender controller 32a connected to three node controllers 34a, 34b, 34c via a data/power bus 30c. The node controller 34a is connected via a data/power bus 30d to an adapter pod 40a, which is in turn connected to a perfusion device 50 in the form of a flow sensor 50a via a bidirectional data/power line 52a. The node controller 34b is connected via a data/power bus 30e to an adapter pod 40b, which is connected to an air embolus sensor 50b via a bidirectional line 52b. The node controller 34c is connected via a data/power bus 30f to an adapter pod 40c, which is connected to a blood pump 50c via a bidirectional line 52c.

[0032] The network extender 22b includes an extender controller 32b connected to three node controllers 34d, 34e, 34f via a data/power bus 30g. The node controller 34d is connected via a data/power bus 30h to an adapter pod 40d, which is connected to a pressure sensor 50d via a bidirectional line 52d. The node controller 34e is connected via a data/power bus 30i to an adapter pod 40e, which is connected to a temperature sensor 50e via a bidirectional line 52e. The node

controller 34f is connected via a data/power bus 30j to an adapter pod 40f, which is connected to a flow occluder 50f via a bidirectional line 52f.

[0033] The main controller 20 is operatively coupled to a blood pump 50g via a bidirectional line 52g connected to an adapter pod 40g. The pod 40g is 5 connected to the main controller 20 via a data/power bus 30k. The main controller 20 is operatively coupled to a level sensor 50h via a bidirectional line 52h connected to an adapter pod 40h, which is connected to the main controller 20 via a data/power bus 30l.

[0034] As used herein, the term "perfusion device" is a device designed to 10 be used in a medical perfusion system, including but not limited to a blood pump such as a centrifugal or roller pump, a flow sensor, a pressure sensor, a temperature sensor, a level sensor, an air embolus sensor or an occluder.

Mechanical Structure of Network Components

[0035] FIG. 2 is a perspective view of a portion of one mechanical embodiment of the main controller 20. Referring to FIG. 2, the main controller 20 has four network connectors 60, which are shown schematically. Each of the 15 network connectors 60 is identical and has the same connector configuration. FIG. 5 illustrates the structure of the connectors 60. As shown in FIG. 5, each connector 60 may be, for example, a standard personal computer connector having nine conductive pins 62 partially surrounded by an asymmetrical metal housing 64.

[0036] FIG. 3 is a perspective view of one embodiment of the network extenders 22 shown schematically in FIG. 1. Each network extender 22 has a hexahedral housing 66 with one side 68 on which three connectors 70 are disposed and an opposite side on which a connector 72 is disposed. Each connector 70 is 25 identical to the connectors 60 and has the structure shown in FIG. 5. The connector 72, which is shown in FIG. 6, has nine pin receptacles 74 formed in an asymmetrical housing 76 composed of an insulating material such as plastic. The pin receptacles 74 are located to correspond to the positions of the nine pins 62 of

the connector 60. Consequently, the connector 72 has the same connector configuration as the connector 60 and thus can be plugged into the connector 60.

[0037] FIG. 4 is a perspective view of the adapter pods 40 shown schematically in FIG. 1. Referring to FIG. 4, each adapter pod 40 has a hexahedral housing with one side 82 on which a connector 84 is disposed and an opposite side on which a connector 86 is disposed. The connector 86 is identical to the connectors 72 described above (and shown in FIG. 6).

[0038] The connector 84 is adapted to be connected to a device connector (not shown) that is associated with one of the perfusion devices 50 described above. The connector 84 has a different connector configuration than the connectors 60, 70, 72, 86. One example of the structure of the connector 84 is shown in FIG. 7 to include six conductive pins 88. Since each of the adapter pods 40 is adapted to be connected to a different type of perfusion device 50 (the pumps 50c, 50g may be different types of pumps, such as a roller pump or a centrifugal pump), the connector 84 disposed on each of the adapter pods 40 may have a different connector configuration.

[0039] Since the connectors 60 of the main controller 20 and the connectors 70 of the network extenders 22 have the same connector configuration as the connector 86 of the adapter pods 40, it should be noted that any of the adapter pods 40 may be plugged into any of the connectors 60, 70. As a result, any combination of perfusion devices 50 may be connected to the main controller 20.

[0040] FIG. 8 illustrates the main controller 20 having the network extenders 22 and the adapter pods 40 connected to it. Each of the adapter pods 40 of FIG. 8 would be connected to a respective one of the perfusion devices 50 shown in FIG. 1 via a respective connector (not shown) attached to the perfusion device 50 by a cable.

[0041] Although the form of the network extenders 22 shown in FIGS. 3 and 8 makes the resulting control unit compact, network extenders having different structures could be used. For example, instead of having the connector 72 fixed on

the housing 66, the connector 72 could be connected to the housing 66 via a cable. Alternatively, the housing 66 could be eliminated, and the connectors 70, 72 could be interconnected via cables.

Electronics

5 [0042] FIG. 9 is a block diagram of the main controller 20 shown schematically in FIG. 1. Referring to FIG. 9, the main controller 20 has a microprocessor (MP) 100, a random-access memory (RAM) 102, a nonvolatile memory 104 such as a hard disk or a flash RAM, a network controller 106, a drawing controller 108, and an input/output (I/O) circuit 110, all of which are interconnected by an address/data bus 112. The I/O circuit 110 is connected to a display device 114, such as a CRT or a flat-panel display, and an input device 116, such as a keyboard or electronic mouse or a touch screen on the display device 114.

10 [0043] The main controller 20 also includes a power supply circuit 118 that is connected to an outside source of AC power and which includes an internal transformer (not shown) that generates +5 volt and +24 volt DC power on a pair of electrical power lines relative to a ground line, which lines are schematically designated 120 in FIG. 9. The electrical power and ground lines 120 are provided to each of four node controllers 34g-34j via a data/power bus 30m and to the other node controllers 34 via the other portions of the network bus 30. The data/power bus 30m includes a number of data communication lines which are connected to the network controller 106.

15 [0044] FIG. 10 is a block diagram of the extender controller 32a shown schematically in FIG. 1 (the design of the extender controllers 32a, 32b is the same). Referring to FIG. 10, the extender controller 32a has a controller 130 and a switch 132, both of which are connected to the data/power bus 30a. The extender controller 32a is connected to its parent node controller 34g via a bidirectional signal line 133. As used herein, a "parent" device is a connected device that is closer to the network controller 106 (FIG. 9) of the main controller 20. The node

controller 34g transmits a unique physical address to the extender controller 32a via the line 133, and the extender controller 32a includes a driver circuit 135 which is used to periodically transmit a check-in code to the node controller 34g via the line 133. The check-in code and the physical address may be the same

5 binary code.

[0045] FIG. 11 illustrates a block diagram of the node controller 34a shown schematically in FIG. 1 (the design of all the node controllers 34 is the same). Referring to FIG. 11, the node controller 34a has a controller 140 which receives an enable signal or a disable signal from the extender controller 32a via one of the lines 134 and a periodic check-in code from the adapter pod 40a via the line 133. The controller 140 is connected to a code generator 144 via a multi-signal line 146. The code generator 144 generates a predetermined multi-bit binary code that uniquely specifies the physical address of the node controller 34a. The code generator 144 may be, for example, a number of printed metal circuit lines, one line for each bit of the code, each line being selectively connected either to +5 volts (logic "1") or to ground (logic "0").

[0046] The controller 140 selectively operates a switch 150 that either connects or disconnects a data bus 152, which may be composed of two individual data lines, that is part of the data/power buses 30c, 30d (and the other buses 30 that make up the network). When the switch 150 is open, the data buses 30c, 30d are disconnected, and when the switch 150 is closed, the buses 30c, 30d are connected to enable data communications between the adapter pod 40a and the other devices connected to the network 30.

[0047] The controller 140 also operates a switch 154 that controls whether +24 volt DC power (relative to a ground line 120c) on a electrical power line 120a is supplied to the adapter pod 40a and a switch 158 that controls whether +5 volt DC power on an electrical power line 120b is supplied to the adapter pod 40a. The electrical power lines 120a, 120b are part of the data/power buses 30c, 30d and the other buses 30 that make up the network. A resistor 162 is connected in

parallel with the switch 154, and a resistor 164 is connected in parallel with the switch 158. The resistors 162, 164 act as current-limiting resistors which prevent large amounts of current from being drawn from the power lines 120a, 120b when the switches 154, 158 are open. The controller 140 is connected to a driver circuit 170 which is used to transmit the physical address generated by the code generator 144 to the adapter pod 40a via the line 133.

[0048] FIG. 12 illustrates a block diagram of the adapter pod 40a shown schematically in FIG. 1. Referring to FIG. 12, the adapter pod 40a has a controller 180 which is powered by a power supply 182 connected to the electrical power lines 120a, 120b. The controller 180 may transmit a check-in code on the line 133 via a driver 184. The controller 180 receives network messages from the data bus 152 and transmits messages onto the data bus 152 via a transceiver 186.

[0049] The controller 180 is connected to a memory 188 and to a device interface circuit 190. The device interface circuit 190 has a plurality of data lines 192 and a plurality of electrical power lines 194 which are connected to the perfusion device 50a via the connector 84 (FIG. 7). The controller 180 causes various types of data signals to be transmitted to the perfusion device 50a via the data lines 192.

[0050] Depending on the type of perfusion device 50 to which an adapter pod 40 is connected, the signals on the data lines 192 might include, for example, digital or analog signals (e.g. 4-20 ma signals) relating to the control of the perfusion device 50, such as a desired pump speed or mode of operation. The number of data lines 192 used depends on the particular perfusion device 50 to which the adapter pod 40 is connected.

[0051] The controller 180 also causes various types of electrical power to be transmitted to the perfusion device 50 via the power lines 194. These types of power include, for example, +5 volt DC power or +24 volt DC power. If power of another voltage level is necessary, the power supply circuit 182 may comprise a DC/DC converter.

Configuration and Display of Perfusion Circuit

[0052] Prior to using the perfusion system 10 for a medical procedure, the operator connects the desired perfusion devices 50 to the main controller 20 by physically connecting the desired adapter pods 40 and/or network extenders 22 to the main controller 20, as shown in FIG. 8.

[0053] Prior to the commencement of a medical procedure, the perfusion system 10 is configured during a configuration process illustrated in FIG. 13A, which is a flowchart of a configuration computer program routine 200 executed by the main controller 20. Referring to FIG. 13A, at step 202 the program generates a visual prompt to the operator to request whether a previous configuration file should be loaded from the memory 104 of the main controller 20. A configuration file generally includes image data corresponding to an image of a perfusion circuit, which may include an outline of the patient, images of a plurality of fluid conduits connected to the patient, and images of the various perfusion devices 50 used in the system 10. Each perfusion device 50 may be represented by a different image, depending upon the type of perfusion device. For example, pumps may be represented by a pump image, whereas a flow sensor may have a different image.

[0054] The configuration file may also include data relating to the perfusion devices 50, such as the manufacturer and model number of the device, the desired operational mode of the device, numeric limits at which an alarm should be triggered, and identification of any associated perfusion device. Two perfusion devices may be "associated" if one device that is used to control a physical process, referred to herein as a control device, is to receive feedback from another perfusion device, referred to herein as a sensing device.

[0055] For example, referring to FIG. 1, the pump 50g could be controlled based on feedback generated by either the level sensor 50h (which would generate a signal indicative of fluid level within a fluid reservoir) or the flow sensor 50a. In the former case, the pump 50g could be controlled to maintain a predetermined level of fluid within the reservoir, and in the latter case the pump 50g could be

controlled to maintain a predetermined flow through the conduit. Any type of conventional feedback control could be used, such as proportional-integral (PI) or proportional-integral-derivative (PID) control. Where it is desired to control the pump 50g based on the output of the level sensor 50h, the association of the pump 50g with the level sensor 50h would be stored in the configuration file.

- [0056] Referring to FIG. 13A, if the operator requested the loading of a configuration file at step 202, the program branches to step 204 where the operator is prompted to select one of those configuration files. If the operator did not want to retrieve a previously stored configuration file, the program branches to step 10 206, where the operator selects one of a predetermined number of types of perfusion circuit images. Each perfusion circuit image could correspond to the perfusion circuit that would be utilized for a different medical procedure. Two different types of perfusion circuit images are illustrated in FIGS. 14A, 14B described below.
- 15 [0057] At step 208, either the perfusion circuit image corresponding to the configuration file selected at step 204 or the perfusion circuit image selected at step 206 is displayed on the display 114. A pair of exemplary perfusion circuit images that may be displayed on the display 114 are illustrated in FIGS. 14A and 14B.
- 20 [0058] Referring to FIG. 14A, an image 232 of a perfusion circuit corresponding to a left ventricle assist device (LVAD)) configuration is shown. The perfusion circuit image 232 includes a patient image 234, an image 236 of a fluid conduit which removes blood from the left ventricle of the patient, an image 238 of a pump, an image 240 of a fluid conduit which returns blood to the aorta of 25 the patient, an image 242 of a flow occluder, an image 244 of a temperature sensor, and a pair of images 246 of an air embolus sensor.
- [0059] Referring to FIG. 14B, an image 248 of a perfusion circuit corresponding to a bi-ventricular assist device (Bi-VAD) configuration is shown. The perfusion circuit image 248 includes all of the images shown in FIG. 14A, as

well as an image 250 of a second blood pump, an image 252 of a conduit which removes blood from right ventricle of the patient, and an image 254 of a conduit that returns blood to the pulmonary artery of the patient. Each of the perfusion circuit images illustrated in FIGS. 14A, 14B could be pre-stored in the memory 104 of the main controller 20, in addition to other types of perfusion circuit images.

[0060] At step 210, the operator may select one of a number of configuration options to change the configuration of the perfusion system 10. If the operator selects the option of adding a perfusion device 50 as determined at step 212, the program branches to step 214 where the operator is prompted to select a type of perfusion device 50, such as a pump or a flow sensor, to add to the perfusion circuit image displayed on the display 114. At step 216, the operator selects the position at which an image of the newly selected perfusion device 50 will be displayed. This position could be specified by the operator via an electronic mouse, and the displayed perfusion circuit image could include a number of possible connection points 256 (FIG. 14A) at which the perfusion device 50 could be connected. The possible connection points 256 could be highlighted, such as by placing them in a bold color or making them blink on and off, so that the possible connection points 256 are readily apparent to the operator. After the operator selects the position, at step 218, an image of the perfusion device is displayed in the perfusion circuit image at that position.

[0061] If the operator selected the option of configuring one of the perfusion devices as determined at step 220, the program branches to step 222 where the current configuration of the perfusion device 50 is displayed next to the image of the device in the perfusion circuit. As noted above, the current configuration could include the mode of operation of the perfusion device, alarm limits for the device, any associated perfusion devices, etc. At step 224, the operator may change or add to the current configuration.

- [0062] If the operator selected the option of displaying data for the perfusion devices as determined at step 226, the program branches to step 228 where it checks to determine whether there is data available to display. Such data could include, for example, the manufacturers and model numbers of the perfusion devices. If there is data available as determined at step 228, the program branches to step 230 where the data is displayed next to the perfusion devices in the perfusion circuit.

Connecting Perfusion Devices

- [0063] The main controller 20 may utilize a plug-in procedure to accommodate perfusion devices 50 that are subsequently connected to the perfusion system 10. FIG. 13B is a flowchart of a plug-in routine 260 performed by the main controller 20. During the plug-in routine, the main controller 20 may operate in an automatic match mode in which it can match a previously entered device configuration with a perfusion device that is subsequently connected to the main controller 20. For example, a operator may configure a centrifugal blood pump (not yet connected to the main controller 20) to operate in a continuous mode to continuously pump a predetermined flow. When the blood pump is subsequently connected to the controller 20, the controller 20 will then automatically match the previously stored pump configuration with the pump.
- [0064] Referring to FIG. 13B, at step 262, if the main controller 20 is in the automatic match mode, the program branches to step 264 where it determines whether there is only one possible match between the perfusion device just connected and the previously stored device configurations. This would be the case where there is only one previously stored configuration for a pump and where the device that was just connected to the main controller 20 was a pump.
- [0065] If there was only one possible match as determined at step 264, the program branches to step 266 where it determines whether the position at which the device is to be displayed in the perfusion circuit image is known. This position could be included in the previously stored configuration for the device. If the

position is not known, the program branches to step 268 where the operator is prompted to select a position, and then the program branches to step 270 where an image of the newly connected perfusion device is displayed in the perfusion circuit image. If the position of the device as determined at step 266 was known, the

5 program skips step 268 and branches directly to step 270.

[0066] If the main controller 20 was not in the automatic match mode, as determined at step 262, or if there was more than one possible match, as determined at step 264, the program branches to step 272. If the newly connected device has already been configured, the program branches to step 274 where the

10 operator is prompted to select the proper configuration from a plurality of prestored configurations. If the device is not already configured, the program branches to step 276 where the operator enters the desired configuration parameters for the device. The program then performs steps 268 and 270 described above.

15 Network Communication and Message Protocol

[0067] The network controller 106, shown in FIG. 9, oversees the flow of data on the network busses 30, each of which includes the data bus 152 (which may be composed of two wires) on which digital data packets are transmitted and received. Each of these digital data packets contains a message. Messages convey

20 data and other information between the various devices (e.g., the main controller 20, the extender controllers 32 and the perfusion device adaptor pods 40) connected to the perfusion system network.

[0068] In the context of a communication network, a protocol is a set of rules describing how data is to be transmitted over the network. The protocol may

25 define, for example, certain electrical and physical standards associated with the transmission of data. It may also define data formatting (e.g., the syntax of the message carrying the data) and, among other things, message priority. In accordance with a preferred embodiment of the present invention, the communication network and messaging protocol (herein "protocol") associated

with perfusion system 10 provides a data format for the data packets as shown in FIG. 18.

[0069] As shown in FIG. 18, each data packet includes a number of data fields: a start-of-frame (SOF) field; an arbitration field; a control field; a variable length data field; an error detection/correction field, such as a cyclic-redundancy-check (CRC) field; an acknowledgment (ACK) field and an end-of-frame (EOF) field. The SOF is, for example, a 1-bit field that marks the beginning of the data packet, which follows an inter-message or inter-data packet idle period. The arbitration field defines, among other things, the type of message that is being conveyed in the data packet. The arbitration field is of particular importance; as such, it will be discussed in much greater detail below. The control field defines the length or, more generally, the amount of data (e.g., the number of data bytes) contained in the variable-length data field that follows. The error detection/correction field, as the name suggests, provides the ability to identify certain bit errors, as well as the ability to correct them. The ACK field is also a 1-bit field which, when properly set, indicates that a receiving device has accurately received and/or processed a corresponding, previously transmitted message. Finally, the EOF field marks the end of the data packet. It will be understood that the format shown in FIG. 18 is exemplary, and that other formats may be employed without departing from the spirit of the present invention.

[0070] As stated, the arbitration field is of particular importance. In accordance with the data packet format illustrated in FIG. 18, the arbitration field has a length of 29 bits. However, the protocol could provide an arbitration field having more than or fewer than 29 bits. The value of these 29 bits define, among other things, the category of the message that is being conveyed in the corresponding data packet. The protocol actually defines ten different message categories, though it will be readily apparent that more or fewer than ten messages categories is conceivable. Each of these ten different message categories will be described in greater detail below.

[0071] The protocol also defines a priority schedule for the ten message categories. Consequently, one of the ten message categories has the highest priority. Another one of the ten message categories has the lowest priority, and the remaining eight message categories have a corresponding priority between the highest and lowest priority. One reason for prioritizing the message categories is that it provides a way of determining which one of two messages should be given deference when two devices inadvertently attempt to transmit these messages at the same time. The protocol, and more particularly, the prioritization scheme defined by the protocol gives deference to the message with the higher priority. Thus, the message with the higher priority is transmitted first, whereas the message with the lower priority can be transmitted at a later time.

[0072] The protocol achieves this prioritization scheme by allocating, for example, the 13 most significant bits (MSBs) in the arbitration field for the purpose of defining, among other things, the message category of the message contained in the corresponding data packet. Again, it will be understood that more than or less than 13 bits, or bits other than the 13 MSBs may be used for this purpose. Nevertheless, the 13 MSBs in the arbitration field of a message with the highest priority might have the binary value "0000000000000". The 13 MSBs in the arbitration field of a message with the second highest priority might have the binary value "0000000000001". The 13 MSBs in the arbitration field of a message with the third highest priority might have the binary value of "0000000000010" -- and so on. Consequently, if there is a collision between two messages on the network, the message having the lowest value in the arbitration field will be given deference.

[0073] The protocol also provides a scheme for prioritizing messages which belong to the same message category. In general, the protocol may achieve this by assigning a device priority number to each type of device that may be connected to the perfusion system network. Table I below presents an exemplary list of devices, along with a corresponding device priority number in decimal and

binary form, where a lower device priority number corresponds with a higher priority level. Thus, for example, the main controller 20, which has the device priority number "00000000", has the highest priority level. In contrast, a service pod having the device priority number "01111110" has the lowest priority of the devices listed in Table I. Alternatively, message priority within a given message category may be based on the logical address assigned to each device. This will now be discussed herein below.

	Priority No.	Device	Priority No.	Device
10	0 [0000000]	QCU	41-44 [00101001 - 00101100]	Reserved
	1-4 [00000001 - 00000100]	reserved	45 [00101101]	Temperature Pods
	5 [00000101]	ABD Pods	46-49 [00101110 - 00110001]	Reserved
	6-9 [00000110 - 00001001]	reserved	50 [00110010]	Gas Flow Pods
	10 [00001010]	Level Detector	51-54 [00110011 - 00110110]	Reserved
15	11-14 [00001011 - 00001110]	reserved	55 [00110111]	Knob Pods
	15 [00001111]	Pressure Pods	56-59 [00111000 - 00111011]	Reserved
	16-19 [00010000 - 00010011]	reserved	60 [00111100]	Serial Pods
	20 [00010100]	Flow Meter Pods	61 [00111101]	CD-100 Pods
	21-24 [00010101 - 00011000]	Reserved	62 [00111110]	CDI-500 Pods
20	25 [00011001]	Large Roller Pumps	63 [00111111]	Neptune Pods
	26-29 [00011010 - 00011101]	Reserved	64 [01000000]	Data Transfer Pod
	30 [00011110]	Small Roller Pumps	65 [01000001]	Power Pods
	31-34 [0011111 - 00100010]	Reserved	66-124 [01000010 - 01111100]	reserved
25	35 [00100011]	Centrifugal Pumps	125 [01111101]	LAN I/F Pod
	36-39 [00100100 - 00100111]	Reserved	126 [01111110]	Service Pod
	40 [00101000]	Occluder Pods	127 [01111111]	Unknown Pod

TABLE I

[0074] In general, each device connected to the perfusion system network is assigned a corresponding logical address. Each device also has a corresponding permanent address. A logical or dynamic address is an address that is assigned to a device when the device gains access to the network. A permanent address, also referred to herein as a physical address, is one that is loaded into the device's non-volatile memory at the time of manufacture. In the event where two messages falling within the same message category are simultaneously transmitted from or to separate devices, and where the two separate devices happen to be the same type

of device (i.e., they have the same device priority number as shown, for example, in Table I), the value of the bits in the arbitration field corresponding to the logical or permanent address of the two similar devices may be used to determine message priority. Again, the message containing the lowest logical address value may be given deference over the other.

	Message Category	MSB ₍₂₉₎	Arbitration Field		LSB ₍₁₎
10	A	00000	00000000	cccccccc	cccccccc
	B	00000	00000001	cccccccc	cccccccc
	C	00000	00000010	aaaaaaaa	ssssssss
	D	00000	00000100	cccccccc	ssssssss
	E	00001	cccccccc	cccccccc	dddddddd
	F	00010	cccccccc	cccccccc	dddddddd
	G	00100	kccccccc	cccccccc	ssssssss
	H	10000	cccccccc	dddddddd	dddddddd
	I	10001	cccccccc	ssssssss	ssssssss
	J	11111	11111111	11111111	11111111

Table II

[0075] As stated, the protocol defines ten different message categories. Table II above lists each of these categories in order of priority, with the message category having the highest priority listed first and the message category having the lowest priority listed last. As indicated, each category has a corresponding letter designation A through J. Table II also provides the format of the 29 bit arbitration field associated with each message category. Again, it will be understood that this format (e.g., the number of bits and the arrangement of the bits) is exemplary.

[0076] Turning to category A type messages, category A messages are control messages which are broadcast from the main controller 20 to one, more than one or all other devices connected to the perfusion system network. As shown, the 13 MSBs have the binary value "000000000000000", which identifies

corresponding messages as belonging to category A, which have the highest priority level. The remaining 16 bits, identified in Table II as having the value "cccccccc cccccccc", are reserved for the purpose of defining, for example, a specific type of control message within category A, where the message may pertain to all of the devices connected to the network, a few of the devices connected to the network or a particular one of the devices connected to the network. In the latter case, the 8 least significant bits (LSBs) could be used for the purpose of defining the logical address of the one particular device for which the message is specifically intended. Category A messages include software reset messages. They also include messages which force a particular device off-line. If such a message is directed to one or more devices, each of these devices would be forced into an off-line mode, where their logical addresses are discarded. The main controller 20 might broadcast this type of message if, for some reason, there is not enough power to support all of the devices currently connected to the network.

[0077] Category B messages have the next highest priority. These messages are broadcast by the main controller 20 to one or more extender controllers 32. The 13 MSBs have the binary value "0000000000001"; accordingly, category B messages have a lower priority level than category A messages. The 16 LSBs, which have the value "cccccccc cccccccc", define the specific type of category B message.

[0078] Message category C covers alarm and/or safety messages. The 13 MSBs in the arbitration field of a category C message have the binary value "0000000000010", which identifies the corresponding message as a category C type message having a priority that is lower than that of a category B type message. The 16 LSBs are identified in Table II as having the value "aaaaaaaa ssssssss", where the bits "aaaaaaaa" may be used to define the type of device which generated the alarm and, therefore, the particular alarm or safety condition. The bits "ssssssss" may be used to define the logical address of the device which

generated the message. Category C type messages may be directed to the main controller 20 in order to alert the main controller 20 of a corresponding alarm or safety condition.

[0079] Category C messages may also be used for the purpose of conveying *triggers*. A *trigger* is a message that instructs one or more devices to automatically take some predefined action in response to a particular event, where the trigger source is the device generating the message and the trigger respondent is the device which has been configured to take action in response to that message. For example, an air bubble detection (ABD) pod may detect the presence of air in a corresponding fluid conduit. The ABD pod then generates and broadcasts a category C type message, where the bits "aaaaaaaa" indicate that the message was generated by an ABD pod and that air has been detected, and where the bits "ssssssss" reflect the logical address of the ABD pod which detected the air. In this example, the ABD pod is the trigger source. The respondent device or devices may be a particular blood pump and/or a particular occluder pod. With regard to the occluder pod, the action taken in response to the ABD pod's trigger message may involve clamping-off the corresponding fluid conduit to prevent any additional air from being pumped into the patient's cardiovascular system. With regard to the blood pump (e.g., a roller pump or a centrifugal pump), the response to the ABD pod's message may involve terminating or pausing the pump's actions. In addition to ABD alarm messages, other examples of category C type messages include pressure alarm/alert messages, flow rate alarm/alert messages, temperature alarm/alert messages, fluid level alarm/alert messages, blood pump start and/or stop messages, occluder pod alarm/alert messages and gas mixer pod alarm/alert messages.

[0080] In accordance with a preferred embodiment of the present invention, a category C type message serving as a *trigger* is broadcast over the network by the trigger source. However, only the trigger respondent or respondents are configured to accept, process and respond to this message. The

process of configuring one or more devices to accept, process and respond to a specific message(s) or messages from a specific source device is referred to herein as message-discrimination. Configuring the one or more respondent devices to achieve message-discrimination involves programming an acceptance filter in each 5 of the one or more respondent devices. Message-discrimination and acceptance filters are discussed in greater detail below.

[0081] Category D messages correspond to servo messages, where the term servo refers to a continuous feedback process. Accordingly, servo messages are messages which contain data that is used in a continuous feedback process, where 10 the feedback process involves at least one sensing device (e.g., a flow sensor, a fluid level sensor, a pressure sensor or a temperature sensor) and at least one respondent device (e.g., a blood pump) which has been configured to continuously accept, process and respond to messages generated by the sensing device. The 13 MSBs in a category C type message have the value "0000000000100", which 15 identifies the corresponding message as a category D type message having a priority that is lower than that of a category C type message. The 16 LSBs have the value "cccccccc ssssssss", wherein the bits "cccccccc" may be used to define a particular type of sensing device, and therefore, a particular data parameter (e.g., pressure, temperature, flow rate or fluid level), and wherein the bits "ssssssss" 20 may be used to define the logical address of the sensing device which measured the data contained in the message. In accordance with the data format shown in FIG. 18, the measured data is actually contained in the variable-length data field of the corresponding data packet.

[0082] Previously, it was stated that the pump 50g could be controlled, 25 based on feedback data generated by the flow sensor 50a in order to maintain a predetermined flow through a corresponding fluid conduit. In this example, the flow sensor 50a continuously measures the flow rate in the conduit. The adaptor pod 40a associated with the flow sensor 50a samples these measurements at a relatively high sampling rate. The adaptor pod 40a then generates a category D

type message and inserts therein (i.e., in the variable-length data field) a corresponding data sample. The adaptor pod 40a then broadcasts the message over the network and the pump 50g, which has been configured to accept, process and respond to servo messages associated with the flow sensor 50a, compares the data value contained in the variable-length data field to a set-point. The set-point, which has been previously stored in memory in the pump 50g, represents a desired flow rate. The pump 50g, based on the result of this comparison, then adjusts its output (e.g., increases its output or decreases its output) in order to maintain the condition defined by the set-point. Again, the process of configuring pump 50g to accept, process and respond to messages associated with the flow sensor 50a is called message-discrimination, which will be discussed below.

[0083] Category E messages correspond to general messages broadcast by the main controller 20 to a specific logical address. In a category E type message, the 13 MSBs have the value "00001 cccccccc", where the 5 MSBs specifically identify the corresponding message as a category E type message having a priority that is lower than category D type messages. As indicated in Table II, bits 9-24 have the value "ccccccc ccccccc". These bits define, for example, the specific type of general message that is being conveyed by the main controller 20. The eight LSBs, identified as "dddddddd", define the logical address of the device for which the message is intended.

[0084] Category E type messages provide the main controller 20 with the ability to dynamically control or configure a specific device. For example, it was previously stated that a particular device can be configured to accept, process and respond to category C type messages from a particular trigger source. Also, as stated, a particular device can be configured to accept, process and respond to category D servo messages from a particular sensor device. Accordingly, the data needed to configure these devices may be conveyed in the variable-length data portion of a category E type message. Examples of category E type messages include: messages which establish a trigger linkage (e.g., a message which

establishes a device as a trigger respondent so that it responds to category C type messages from a particular trigger source); messages which establish a servo linkage (e.g., a message which configures a device so that it responds to category D type servo messages from a particular sensor device); messages which configure a pump to operate at a desired speed; and messages which establish alarm/alert thresholds (e.g., a pressure or flow rate alarm threshold).

[0085] Category F messages are general messages broadcast from a device, other than the main controller 20, to a peer device. These messages allow a device to provide data to, and therefore, exercise some limited control over a peer device.

Category F type messages are particularly important when and if the main controller 20 has failed, or is otherwise unavailable. An example of a category F type message is a message broadcast by a roller pump to clear an ABD alarm in a corresponding ABD pod.

[0086] Table II, once again, provides an exemplary format for the 29 bit arbitration field associated with category F type messages. As indicated, the 13 MSBs, and in particular, the 5 MSBs "00010" identify the message as a category F type message having a priority which is lower than a category E type message. Bits 9-24, which have the value "ccccccc ccccccc" define a specific type of category F message, whereas the eight LSBs "ddddddd" define the logical address of the peer device (i.e., the device which is the intended recipient of the message).

[0087] Category G messages are general messages broadcast from a specific logical address. The 13 MSBs, and in particular, the 5 MSBs "00100" identify the message as a category G type message having a priority which is lower than a category F message. Bits 9-24, which have the value "kcccccc ccccccc" define a specific type of category G message. The eight LSBs "ddddddd" define the logical address of the device which is broadcasting the message.

[0088] There are two general types of category G messages. The first type involves messages broadcast from a particular logical address which convey data,

such as sensor, configuration or status information, at a low data rate as compared to feedback data which is conveyed at a relatively high data rate. Category G messages of this type include messages which convey autonomous current pressure, current temperature and accumulated pump volume.

5 [0089] The second general type of category G message involves broadcasting an acknowledgment to certain previously broadcasted messages. For example, certain category E and category F type messages may require that the recipient device acknowledge that they have accurately received and/or processed the message. If the category G message is, in fact, an acknowledgment message
10 one of the bits in the arbitration field, for example, bit 24 (which is identified by the letter "k" in Table II) may be set to "1" to signal that the message is being sent as an acknowledgment and that the previous message was accurately received and processed. Bits 9-23 of the arbitration field would be identical to bits 9-23 in the previously broadcasted category E or category F type message. Bit 24 would then
15 be set to "0" if the message is not an acknowledgment of a previously broadcasted message.

[0090] Category H messages are general messages which are broadcast to a permanent (i.e., physical) address. These messages are generally used for the purpose of sending data to devices that do not presently have a logical address
20 (e.g., devices that have not yet been assigned a logical address or devices which have been forced off-line). Thus, category H messages may be used to configure logical addresses. They may also be used to perform software upgrades and perform debug or diagnostic operations. Typically, the receiving device is expected to acknowledge receipt of a category H message. This may be
25 accomplished by broadcasting a category I type message -- a broadcast message from a permanent address. Category I messages are described below.

[0091] As shown in Table II, the 13 MSBs of a category H type message have the value "10000 cccccccc", where the 5 MSBs identify the message as a category H message having a priority which is lower than a category G message.

Bits 17-24 define the type of category H message being broadcast, and these bits may also define the type of device for which the message is intended. The remaining bits 1-16, which have the value "ddddddddd dddddd", define the permanent address of the intended recipient device.

5 [0092] Category I messages are general messages broadcast from a particular permanent address. These messages are, in general, used by devices that do not presently have a logical address. More specifically, they may be used to request a logical address or, as stated above, to acknowledge a previously broadcasted category H type message. Table II indicates that the 13 MSBs of a
10 category I type message, and in particular, the 5 MSBs "10001" identify the message as a category I type message having a priority that is lower than a category H message. Bits 17-24 define the specific type of category I message being broadcast and, in some cases, the type of device broadcasting the message.
15 The remaining bits 1-16, which are shown as having the value "ddddd dddddd", define the address of the device broadcasting the message.

[0093] Finally, category J messages correspond to a status request from the main controller 20. Typically, these are messages which are periodically broadcast to all devices connected to the network. Table II indicates that all 29 bits in the arbitration field of a category J message are set to "1". Accordingly,
20 category J messages have the lowest priority.

[0094] As stated above, the data packet format illustrated in FIG. 18 includes a variable-length data field. However, some of the messages described above convey no data in this data field. For example, category J status request messages, which are broadcast by the main controller 20 contain no data in the
25 variable-length data field portion of the data packet. In contrast, category D servo messages, as stated above, contain data in the variable-length data field, where this data specifies the numeric value of a sensed condition, i.e., a corresponding parameter, such as a flow reading of 0.257 liters per minute.

[0095] The main controller 20, the extender controllers 32, and the adapter pods 40 may include conventional electronics for checking the accuracy of received messages via the CRC field, requesting retransmission of messages that were not accurately received, and for transmitting acknowledgment messages in response to the receipt of certain types of messages when they are accurately received.

[0096] The messages described above may be transmitted or broadcast to all the devices connected to the network 30. Each device, such as a pod 40 or an extender controller 32, can discriminate by accepting, processing and, if necessary, responding only to certain broadcasted messages. This ability, as stated above, is referred to herein as message-discrimination. For example, message-discrimination may be accomplished by maintaining an acceptance filter in the various devices. The acceptance filter in a given device may comprise a memory, such as a number of dedicated registers, each of which has the ability to store a logical address of another device in which the given device is interested. In general, these logical addresses are transmitted to the given device during the configuration process by the main controller 20, via a category E type message, as described above.

[0097] For example, if the given device is the adapter pod 40c, which is connected to the blood pump 50c which controls the flow of blood through a conduit based on feedback data from the flow sensor 50a, the acceptance filter in pod 40c would have been programmed during a configuration process to include the logical address of the flow sensor 50a. In doing so, pod 40c would receive, accept, process and, if necessary, respond to any message generated by the pod 40a connected to the flow sensor 50a.

[0098] The acceptance filter would also be programmed to include the logical address of the main controller 20, in addition to the logical addresses of any other devices. It should be noted that message-discrimination precludes the

need to include a specific destination address in a given message (although a destination address may be included).

[0099] The pods 40 and extender controllers 32 could also discriminate messages based on the type of message instead of the identity of the sender. For example, the acceptance filter in a pod 40 could be programmed to receive all category J status-request messages and/or all category E configuration messages. This may be achieved by storing, in a corresponding register, the value of the arbitration field bits which specifically identify these messages as such.

[00100] Before use of the perfusion system 10 for a medical procedure, and after all the perfusion devices 50 are configured as described above, data packets containing configuration messages are transmitted to all the pods 40 connected to the network 30. The configuration messages include all the necessary configuration data described above. For example, the configuration data for a blood pump may include the operational mode of the pump, the desired flow rate of the pump, etc. The configuration data would also include device association data (e.g., the data needed to establish a feedback link between two devices such as a sensing device and a pump, the data needed to establish a trigger link between a trigger source and a trigger respondent, or the data needed to establish a master/slave relationship between two devices, such as two pumps). If, for example, device association data is received by a pod 40, the message-discrimination memory (e.g., the acceptance filter) in the pod 40 would be updated with the logical address of the associated device. It should be noted that these device associations continue to exist and continue to govern the operation of the corresponding devices even when there is a subsequent loss of the main controller 20.

Connecting Pods to the Network

[00101] In order for them to communicate with the main controller 20 via the network data/power buses 30, the adapter pods 40 must be granted permission

to connect to the network 30. This connection is initiated with a startup-request message transmitted to the main controller 20 by the adapter pod 40 for which the network connection is to be made. The startup-request message includes a first code identifying the type of perfusion device 50 connected to the pod 40 requesting to be connected and a second code identifying the physical address (specified by the code generator 144 of FIG. 11) of the pod 40 requesting to be connected.

5 [00102] FIG. 13C is a flowchart of a startup routine 280 that is performed by the main controller 20 in response to the receipt of a startup-request message from an adapter pod 40. The startup routine 280 may be an interrupt service routine which is invoked in response to an interrupt generated upon the receipt of a startup-request message by the main controller 20. Referring to FIG. 13C, at step 10 282, the startup message received by the main controller 20 is decoded to determine the type of the perfusion device 50 attached to the pod 40 requesting to be connected and to determine the physical address of the pod 40.

15 [00103] At step 284, the program determines whether full power should be granted to the requesting pod 40. As described above, electrical power to run the perfusion devices 50 is provided to the pods 40 via the network 30 from a power supply 118 (FIG. 9) in the main controller 20. Since the power available from the power supply 118 may be limited, the main controller 20 may be programmed to 20 allow only a certain number of perfusion devices 50 to be connected to the network 30, or alternatively, to allow only certain numbers of specific types of perfusion devices 50 to be connected. For example, since control devices such blood pumps typically draw more power than sensing devices, the main controller 20 may be provided with an upper limit on the number of control devices that can 25 be connected to the network 30.

[00104] At step 284, the decision whether to grant full power could be made by comparing the number of perfusion devices 50 already connected to the network 30 with the maximum number that can be connected to determine whether the connection of an additional device 50 will cause the maximum to be exceeded.

Alternatively, if the device requesting connection is a control device, then the number of control devices already connected could be compared with the maximum number of control-type perfusion devices that can be connected. If it is determined that full power should not be granted, the program simply ends.

5 [00105] If it is determined that full power should be granted, steps 286-298 are performed to generate and transmit a startup granted message that will cause the pod 40 associated with the perfusion device 50 to be connected to the network 30. In particular, at step 286, a unique logical address is allocated to the newly connected pod 40. For example, where the capacity of the network 30 is sixteen
10 devices, the logical address may be a four-bit binary code. At step 288, a startup-granted message which includes the logical address is encoded, and at step 290 the startup-granted message is transmitted over the network 30.

15 [00106] At step 292, if the pod 40 which requested startup is local to the main controller 20 (i.e. if it is one of the pods 40g or 40h directly connected to the main controller 20 without a network extender 22), full power to the pod 40 is enabled via the local node controller (i.e. one of the node controllers 34i-34j of FIG. 9) connected to the perfusion device 50. Referring to FIG. 11, this is accomplished by sending an enable signal on the line 134, which will cause the controller 140 to close the switches 150, 154, 158 so that full power is supplied on
20 the power lines 120a, 120b and so that the adapter pod 40 is connected to the data bus 152.

25 [00107] Referring to FIG. 13C, if the perfusion device was not a local device as determined at step 292, the program branches to step 296 where a connect message which includes the physical address of the node controller 34 associated with the pod 40 requesting startup is encoded, and then to step 298 where the connect message is transmitted over the network 30. As described below, when the extender controllers 32 receive the connect message, they decode it to determine the physical address, and the extender controller 32 connected to the node controller 34 having that physical address (i.e. the node controller 34

associated with the requesting pod 40) turns on full power by transmitting an enable signal on the line 134 connected to that node controller.

Status Requests

[00108] During operation of the perfusion system 10, to ensure that all devices connected to the network 30 are properly functioning and are receiving messages broadcast over the network 30, the main controller 20 periodically transmits a status-request message to all extender controllers 32 and adapter pods 40 on the network 30. Each extender controller 32 and adapter pod 40 must respond to the status request within a predetermined period of time. Any extender controller 32 or pod 40 that fails to respond to the status request within that time period is disconnected from the network 30, and a corresponding alarm message is generated on the visual display 114 to warn the operator of such event.

[00109] FIG. 13D is a flowchart of a status-request routine 300 periodically performed by the main controller 20. Referring to FIG. 13D, at step 302 a status-request message is encoded, and at step 304 the message is broadcast to all extender controllers 32 and adapter pods 40 connected to the network 30. As described above, the status-request message may simply be all logical "1"s in the arbitration of the data packet. At step 306, a predetermined time-out period within which all extender controllers 32 and pods 40 must respond to the status request message is started.

[00110] Upon receiving the status-request message, each extender controller 32 and adapter pod 40 encodes a status message with its logical address and its status, and then broadcasts the status message to the main controller 20 over the network 30.

[00111] FIG. 13E is a flowchart of a receive status routine 310 that is performed by the main controller 20 upon receipt of a status message transmitted to it in response to the status-request message previously transmitted by the routine 300. Referring to FIG. 13E, at step 312 the logical address of the responding extender controller 32 or adapter pod 40 is determined from the status message,

and at step 314 the status of the device 32 or 40 is determined from the message. The status may be specified by a number of different binary status codes. At step 316, if the status of the device 32 or 40 is okay, the program simply ends.

However, if the status is not okay, the program branches to step 318 where it

- 5 responds to a status condition identified by the status code. If the condition is relatively minor, the main controller 20 may simply generate a warning on the visual display 114 of the perfusion system 10. If the condition is serious enough, the main controller 20 may disconnect the device 32 or 40 from the network 30.

[00112] FIG. 13F is a flowchart of a disconnect routine 330 that causes an extender controller 32 or an adapter pod 40 to be disconnected from the network 30. The disconnect routine 330 is performed by the main controller 20 in response to either: 1) the failure of a device 32 or 40 to transmit a status message to the main controller 20 within the timeout period described above or 2) a serious malfunction of a device 32 or 40 as determined at step 318 of FIG. 13E.

- 10 [00113] Referring to FIG. 13F, at step 332, if the device 32 or 40 to be disconnected is local to the main controller 20, the program branches to step 334 where that device 32 or 40 is disconnected by the node controller 34g, 34h, 34i, or 34j in the main controller 20 that is connected to that device 32 or 40. Referring to FIG. 11, the disconnection is accomplished by transmitting a disable signal on the 15 line 134, which will cause the controller 140 to open the switches 150, 154, 158 so that the data bus 152 and the power lines 120a, 120b are disconnected.

- 20 [00114] At step 332, if the device to be disconnected is not local to the main controller 20, the program branches to step 336 where a disconnect message which includes the physical address of the node controller 34 of the device 32 or 40 to be disconnected is encoded, and then to step 338 where the disconnect message is transmitted over the network 30.

- 25 [00115] If the device to be disconnected is a pod 40 connected to an extender controller 32 via a node controller 34, when that extender controller 32 receives the disconnect message, it decodes it to determine the physical address of

the pod 40 to be disconnected, and the node controller 34 associated with that pod 40 disconnects the pod 40 by transmitting a disable signal on the line 134 connected to that node controller 34.

Operator Commands Input to Main Controller

5 [00116] During operation of the perfusion system 10 during a medical procedure, the main controller 20 responds to various commands and other inputs entered by the operator of the system 10. FIG. 13G is a flowchart of a control command routine 350 which illustrates how the main controller 20 responds to those inputs. While FIG. 13G discloses various possible operator inputs, it should
10 be understood that the main controller 20 could respond to other or additional operator inputs.

15 [00117] Referring to FIG. 13G, at step 352, if the input entered by the operator was a control command, the program branches to step 354 where a control message corresponding to the control command is encoded in a data packet, and then to step 356 where the control message is broadcast over the network 30. For example, the control message could be one of the following: 1) a new alarm limit for a particular sensing device; 2) a new mode of operation for a blood pump; 3) a new target flow value for a blood pump; 4) a new rate at which a particular sensing device should be read; 5) a pump start command; 6) a pump
20 stop command; etc.

25 [00118] At step 358, if the operator requests that a particular alarm be reset, the program branches to step 360 where a corresponding alarm-reset message, which includes the logical address of the device that generated the alarm, is encoded and to step 362 where the alarm-reset message is broadcast over the network 30. At step 364, if the operator requests that the perfusion system 10 be reset, the program branches to step 366 where a corresponding system reset message is encoded and to step 362 where the system reset message is broadcast over the network 30.

Receipt of Network Messages by Main Controller

[00119] During operation, the main controller 20 receives messages of various types that are broadcast over the network 30. FIG. 13H is a flowchart of a receive routine 370 performed by the main controller 20 that illustrates actions taken by the main controller 20 in response to the receipt of various types of messages.

[00120] Referring to FIG. 13H, at step 372, if the received message corresponds to an event message, such as an alarm or other event, the program branches to step 374 where the visual display generated on the display device 114 is updated to advise the operator of that event, and the program branches to step 376 where the event is logged into an event log stored in the memory 104 of the main controller 20.

[00121] At step 378, if the received message corresponds to a data message, such as a message which includes the numeric value representing the output of a flow sensor, the program branches to step 380 where the visual display is updated, and the program branches to step 382 where the data and the device which generated the data are logged into a data log stored in the memory 104 of the main controller 20.

[00122] At step 384, if the received message is a status message, the program branches to step 386 where the status message is processed, as described above in connection with FIG. 13E. At step 388, the visual display is updated based on the status, and at step 390 the status is stored in a status log stored in memory. At step 392, if the received message is a startup request, the program branches to step 394 where the startup request is processed, as described above in connection with FIG. 13C, and at step 396, the visual display is updated.

Operation of Extender Controllers

[00123] A basic function of the extender controllers 32 is to control the connection and disconnection of adapter pods 40 to the network 30. FIG. 15A is a flowchart of a startup routine 420 performed by the controller 130 of each extender controller 32. Referring to FIG. 15A, at step 422 the extender controller

32 performs a number of internal self-tests, such as tests of an internal RAM and an internal ROM. At step 424, if the tests were successful, the program branches to step 426 where the connection of the extender controller 32 to the network bus 30 is tested by transmitting a message onto the network bus 30 and simultaneously receiving the message from the network bus 30 as it is transmitted to determine if the message was in fact transmitted.

5 [00124] At step 428, if the data bus test was successful, the program branches to step 430 where the extender controller 32 starts to periodically transmit a check-in code to its parent node controller 34 via the line 133. As

10 described below, each device (either an adapter pod 40 or an extender controller 32) must periodically transmit a check-in code to its parent node controller 34 to maintain its connection to the network 30.

15 [00125] At step 432, the extender controller 32 waits for its physical address to be transmitted to it from its parent node controller 34. At step 434, if not all of the tests performed at steps 422 and 426 were passed, an error message is broadcast over the network 30. The error message includes the physical address of the extender controller 32 and a binary code which specifies which test(s) were not passed.

20 [00126] If all tests were passed, the program branches to step 438 where a startup-request message containing the physical address of the extender controller 32 is encoded and broadcast over the network 30. At step 440, the program waits until a startup-granted message is received from the main controller 20, and then at

25 step 442 the program waits until full power is granted to the extender controller 32 via the electrical power lines 120a, 120b of its parent node controller 34. When full power is granted, the program branches to step 444 where the extender controller 32 measures the voltages on and the current provided by the electrical power lines 120a, 120b to make sure they are within specification. At step 446, if the power measurements are not within specification, the program branches to step

436 where a message to that effect is broadcast to the main controller 20 over the network 30.

[00127] FIG. 15B is a flowchart of a connect routine 450 performed by an extender controller 32 when it receives a connect message from the main controller 20. Referring to FIG. 15B, at step 452, the connect message received from the main controller 20 is decoded to determine the physical address of the adapter pod 40 to be connected to the network 30. At step 454, the physical address is inspected to determine if the adapter pod 40 to be connected is local to the extender controller 32, meaning that the adapter pod 40 is one of the three that are connected to the extender controller 32. If the pod 40 is not local to the extender controller 32, no further action is taken and the routine 450 ends. If the pod 40 is local to the extender controller 32, the program branches to step 456 where an enable signal is transmitted via one of the lines 134 to the node controller 34 associated with the adapter pod 40 to be connected, which causes the adapter pod 40 to be connected to the network 30 in the manner described above.

[00128] When an extender controller 32 receives a disconnect message from the main controller 20, a disconnect routine 460 shown in FIG. 15C is performed by the extender controller 32. Referring to FIG. 15C, at step 462, the disconnect message received from the main controller 20 is decoded to determine the physical address of the adapter pod 40 to be disconnected to the network 30. At step 464, the physical address is inspected to determine if the adapter pod 40 to be disconnected is local to the extender controller 32. If the pod 40 is not local, no further action is taken. If the pod 40 is local, the program branches to step 466 where a disable signal is transmitted via one of the lines 134 to the node controller 34 associated with the adapter pod 40 to be disconnected, which causes the adapter pod 40 to be disconnected to the network 30.

Operation of Node Controllers

[00129] The basic function of the node controllers 34 is to connect and disconnect the adapter pods 40 (if a node controller 34 is the parent of an adapter

pod 40) and the extender controllers 32 (if a node controller 34 is the parent of an extender controller 32) from the network 30. The connection or disconnection is performed pursuant to an enable or disable signal received either from the main controller 20 or from the extender controller 32 associated with the node controller 34, as described above. In addition, each node controller 34 requires its associated device 32 or 40 to periodically check-in. If the device 32 or 40 fails to check in with a proper check-in code, the node controller 34 disconnects the device 32 or 40 from the network 30.

[00130] FIG. 16A is a flowchart of a node routine 470 performed by each of the node controllers 34. The routine 470 is performed upon the receipt by the node controller 34 of a check-in code received from the associated device 32 or 40.

Referring to FIG. 16A, at step 472, if the code received from the device 32 or 40 is not valid, no further action is taken and the routine ends. The check-in code may be the physical address specified by the code generator 144 (FIG. 11) of the node controller 34. To determine whether the code is valid, the node controller 34 may compare the received code to determine whether it matches a predetermined code.

[00131] If the check-in code was valid, the program branches to step 474 where a time-out timer is restarted. The time-out timer tracks the predetermined period of time within which the device 32 or 40 must transmit a valid check-in code. At step 476, the node controller 34 transmits the physical address generated by the code generator 144 to the pod 40. At step 478, the node controller 34 connects the device 32 or 40 to the data bus 152 (FIG. 11) by sending a signal to the switch 150 which causes it to close (or remain closed if it was already closed).

[00132] At step 480, if the enable signal is present on the line 134 connected to the node controller 34, the node controller 34 supplies full power to the device 32 or 40 by sending signals to the switches 154, 158 (FIG. 11) to cause them to close (or remain closed if they were already closed).

[00133] If the enable signal was not present as determined at step 480, the program branches to step 484, where the device 32 or 40 is disconnected from the

data bus 152 by opening the switch 150 and to step 484 where the device 32 or 40 is disconnected from the electrical power lines 120a, 120b by opening the switches 154, 158.

[00134] If the device 32 or 40 fails to transmit a valid check-in code to the node controller 34 within the time-out period, a time-out routine 490 shown in FIG. 16B is performed by the node controller 34. Referring to FIG. 16B, at step 492 the device 32 or 40 is disconnected from the data bus 152 by opening the switch 150, and at step 494 the device 32 or 40 is disconnected from the electrical power lines 120a, 120b by opening the switches 154, 158.

Operation of Adapter Pods

[00135] The adapter pods 40 perform a number of functions, including receiving configuration and control messages transmitted by the main controller 20, receiving sensing messages containing numeric values of sensed conditions, such as flow, and/or transmitting sensing messages over the network 30. These functions are described below.

[00136] FIG. 17A is a flowchart of a startup routine 520 performed by the controller 180 of each adapter pod 40. Referring to FIG. 17A, at step 522 the adapter pod 40 performs a number of internal self-tests, such as tests of an internal RAM and an internal ROM. At step 524, if the tests were successful, the program branches to step 526 where the connection of the pod 40 to the data bus 152 is tested by transmitting a message onto the data bus 152 and simultaneously receiving the message from the data bus 152 as it is transmitted to determine if the message was in fact transmitted.

[00137] At step 528, if the data bus test was successful, the program branches to step 530 where the adapter pod 40 starts to periodically transmit a check-in code to its parent node controller 34 via the line 133. At step 532, the pod 40 waits for its physical address to be transmitted to it from its parent node controller 34. At step 534, if not all of the tests performed at steps 522 and 526 were passed, an error message is broadcast over the network 30. The error

message includes the physical address of the pod 40 and a binary code which specifies which test(s) were not passed.

[00138] If all tests were passed, the program branches to step 538 where a startup-request message containing the physical address of the adapter pod 40 is encoded and broadcast over the network 30. At step 540, the program waits until a startup-granted message is received from the main controller 20, and then at step 542 the program waits until full power is granted to the adapter pod 40 via the electrical power lines 120a, 120b of its parent node controller 34. When full power is granted, the program branches to step 544 where the pod 40 measures the voltages on and the current provided by the electrical power lines 120a, 120b to make sure they are within specification. At step 546, if the power measurements are not within specification, the program branches to step 536 where a message to that effect is broadcast to the main controller 20 over the network 30.

[00139] During operation, an adapter pod 40 may receive control or configuration messages from the main controller 20 over the network. FIG. 17B is a flowchart of a receive routine 550 that is performed when the adapter pod 40 receives a message. Referring to FIG. 17B, at step 552 the message is decoded to determine the control command embedded in the message, and at step 554, the pod 40 transmits a control signal (via one or more of the data lines 192 in FIG. 12) to the perfusion device 50 connected to it.

[00140] During operation, an adapter pod 40 may receive an alarm signal from the perfusion device 50 via one of the data lines 192. When such an alarm signal is received, an alarm routine 560 shown in FIG. 17C is performed by the pod 40. Referring to FIG. 17C, at step 562 an alarm message is encoded with the logical address of the perfusion device 50 that generated the alarm and the type of alarm, and at step 564 the alarm message is broadcast over the network 30 to the main controller 20.

[00141] During operation, each adapter pod 40 connected to a perfusion device 50, such as a flow sensor, which generates a sensing signal periodically

reads the numeric value of the sensing signal via one of the lines 152. The time period between successive readings of the sensing signal may be specified during the configuration process as described above. FIG. 17D is a flowchart of a sensing routine 570 that is performed when it is time to read the value of the sensing

5 signal. Referring to FIG. 17D, at step 572 the sensing signal is read via one of the data lines 152. At step 574, the numeric value of the sensing signal is encoded in a message along with the logical address of the perfusion device 50 which generated the sensing signal. At step 576, that message is then broadcast over the network 30 to all devices connected to the network 30.

10 [00142] As described above, each adapter pod 40 connected to the network 30 may be provided with a message-discrimination circuit which is used to selectively receive messages from only a subset of the devices connected to the network 30. When the sensing message is broadcast at step 576, the only devices that receive it are the main controller 20 (which may receive all messages broadcast over the network 30) and the particular perfusion device 50 which is being controlled based on the value of the sensing signal encoded in the sensing message.

15 [00143] It should be understood that the adapter pods may be provided with additional functionality not described above. Also, instead of having electrical power being distributed over the network from a single power source provided in the main controller 20, electrical power could be distributed from a plurality of power sources, for example, from one power source provided in each of the network extenders.

20 [00144] Since numerous additional modifications and alternative embodiments of the invention will be apparent to those skilled in the art in view of the foregoing description, the above description is to be construed as illustrative only, and is for the purpose of teaching those skilled in the art the best mode of carrying out the invention. The details of the structure may be varied substantially

without departing from the spirit of the invention, and the exclusive use of all modifications which come within the scope of the appended claims is reserved.